



BANK OF ENGLAND

Speech

Covid and the composition of spending

Speech given by

Ben Broadbent, Deputy Governor, Monetary Policy

Bank of England webinar

12 January 2021

I've received helpful comments from colleagues at the Bank of England. I'd like to thank Cristiano Cantore, Rupert de Vincent-Humphreys, Pavandeep Dhani, Federico Di Pace, Riccardo Masolo, Silvia Miranda-Agrippino, Mette Nielsen and Carleton Webb for their help in preparing the speech. The views expressed are my own and do not necessarily reflect those of the Bank of England or other members of the Financial Policy Committee or the Monetary Policy Committee.

Introduction

Good morning. Thank you for making the time to attend this talk.

It is dispiriting to have to give it as we enter another national lockdown. But the mass vaccination programme now underway also provides genuine hope that this latest set of measures may be the last. And in the meantime, as far as the economic effects of the pandemic are concerned, we should take the opportunity to learn from what happened in 2020. In this respect I want to focus today on one particular aspect of the experience last year, namely the huge shifts in the pattern of consumer spending.

Chart 1 is one place to start. It plots annual growth in GDP, from one fourth quarter to the next, against the same for retail sales¹. As you all know, the first lockdown led to an enormous drop in economic activity, of over 20%, during the first half of 2020. And despite a very sharp recovery in the summer, as the first lockdown was lifted, it's likely that GDP in the fourth quarter of last year was around 10% lower than at the end of 2019. That would make it the sharpest decline, through any calendar year, at least since 1920 (as far back as quarterly estimates go).

Yet, perhaps befitting a year of extraordinary events, last year's downturn was very unusual not just in its size and speed, but in its shape as well. If you had only this correlation to go by, and knew what had happened to GDP, you'd also expect to have seen a steep decline in retail sales. The two series don't always move in lockstep but, as you can see, they've been reasonably well correlated over the past. If that's been a rule, however, then what we saw in 2020 was a very clear exception to it: GDP growth may have been the weakest on record but retail spending growth is just about the strongest.

Chart 1: Weak GDP growth, strong retail sales



Sources: ONS and Bank calculations.
Retail sales data have been adjusted for the switch to CPI deflation in 2001. The figure for 2020Q4 is based on data to November, GDP growth on the MPC's November projection.

Chart 2: Big divergence across consumption



Sources: ONS and Bank calculations.
The retail sales data have been adjusted for the switch to CPI deflation in 2001. Retail sales growth for 2020 Q4 is based on data to November, services growth on data available to date.

¹ The series is retail sales excluding fuel, which goes back slightly further (to 1988) than the aggregate figure.

There are probably several reasons for this anomaly, and in the first part of this talk I'll briefly go through some of them. Perhaps the single most important factor is the introduction of the various furlough schemes. At the cost of a huge jump in government debt these have shielded household income from the drop in national income seen in the first half of last year. Spending by today's consumers has in part been sustained and funded by tomorrow's taxpayers.

However, what I want to spend more time on in this talk is something else: the significant shifts across different areas of that spending. Retail sales may have been strong – they've grown by more than the historical average and more than household income – but, as you can see from Chart 2, spending on consumer services has been extremely weak.

There's an obvious reason for this, of course: in the midst of the pandemic some activities involve more risk of infection than others. Going to a busy shop may be more risky than shopping online. Sadly, eating in a restaurant, going to a football game and attending a concert are riskier still. Whether because of people's own natural caution, or the official lockdowns, spending on anything that involves physical proximity to other people has had to be cut significantly.

But cutting back on these activities doesn't necessarily mean a similar decline in total consumption. At least in principle, people could compensate for cuts in some areas by spending more in others – switching from physical to online retail, for example, or from restaurants to take-away food. And my aim today is to try and understand a bit about the extent of these shifts in consumer demand and their consequences for the supply side of the economy.

As far as the demand side is concerned, a relevant question is this: faced with a rise in infection risk, and a necessary curtailment in the activities that expose people to it, will people choose to spend all the money they would've done, simply transferring it to other, less risky things? Or would you expect them instead to reduce their total spending and save a little more, essentially deferring consumption to the future?

To help think about this question I'll use the simplest and most stylised model of consumer behaviour imaginable. It focuses only on the direct effects of the pandemic on consumption and saving and ignores everything else. Obviously this isn't realistic. For example, we know that saving rates tend to rise when people become more pessimistic or more uncertain about their future income. It's clear that the crisis could have indirect effects of this sort.

But I think it's helpful to separate the differing roles of economic and pandemic risk on demand. And as far as the latter is concerned, it turns out that two things matter for this choice (i.e. whether to switch to other things the money people would've been spending on risky activities or instead to save it). One is how similar the two sorts of consumption are. If you don't lose much by substituting one for the other then, in the face of

a jump in infection rates, you're likely to do more of it. The compensating rise in non-risky spending would be bigger and the decline in total spending smaller.

The other factor is the expected duration of the jump. At least as predicted by this simple model, the degree of expenditure switching is greater, and the drop in aggregate spending smaller, the longer the rise in infection rates is expected to persist. It's when the opposite is true – when people expect pandemic risk to dissipate and conditions for spending to improve relatively soon – that they're more inclined to wait, and defer some spending (even in aggregate) until that happens. So at least as far as its direct impact is concerned, we might expect the pandemic to have material effects on aggregate consumption only if those effects are also temporary. Simple theory suggests there should be an inverse relationship between the depth of the downturn and its duration. Seen this way, the severity of the drop in consumption last spring was in some ways the counterpart to its brevity.

The second thing I want to touch on today is what such expenditure-switching does to the supply side and average costs in the economy. The answer for relatively moderate shifts is probably “not much”. The composition of what we buy changes all the time, as do relative prices, and we can usually safely ignore this when thinking about the economy as a whole. Nor should such shifts in the composition of spending matter much in the longer term, even if they're sizeable. That's because, over time, economic resources – people, capital and general know-how – can move in response. The economy adapts.

But for a short time, at least, significant shifts in demand may be harder to absorb. They result in excess capacity in some areas and too little in others. Until they're able to respond to the shift, resources are therefore misallocated. As I'll explain later on, this can have negative effects on productivity, and positive effects on average costs, even for a given level of aggregate demand. The consequences for cost pressures over the medium term – and therefore for monetary policy – are probably limited. But this divergence in demand may help to explain why, at least on impact, the pandemic has depressed inflation by a bit less than we and others anticipated when it began.

Various things have contributed to gap between real GDP and retail spending


Expenditure switching certainly isn't the only reason for the contrast in Chart 1. I want first to run through a number of other things that have probably contributed to it.

One is the weakness of spending – on durables in particular – during the lockdown in the spring. At that time there were many things people were simply unable to buy. For durable stuff, much of that spending is likely to have been delayed rather than cancelled. If you needed a new fridge or a new car in March, you presumably still needed one in June, when the lockdown began to be eased. One would've expected most of that catch-up to occur during the summer but some of it may also have leaked into the autumn and early winter.

Second, part of the drop in GDP reflects an unusually large decline in the measured real output of public services, something that has no consequences for people’s incomes or their ability to spend. To construct these estimates the Office for National Statistics uses detailed information on things like the numbers of pupils in school and operations in hospitals. Normally these things carry on uninterrupted during recessions: output of public services is in general less volatile than private-sector output (Chart 3). But on this occasion, sweeping all before it, the pandemic curtailed many of these more standard activities, even after the full lockdown was ended in the summer. In the third quarter of this year, for example, real healthcare output – as estimated by the ONS – was still more than 20% lower than in 2019Q3. Taken together, measured declines in measured health and education output accounted for a quarter of the near-9% drop in aggregate GDP during the course of that year.


Note that there’s been no decline in nominal spending on public services, or the collective incomes of people who work in them. Employment in public services has continued to rise and so have average wages. It’s just that, at least as measured in the national accounts, productivity in that sector has declined and the implicit price of those services – the so-called “deflator” – has risen. The result is an additional wedge, relative to any normal downturn, between (measured) real GDP and household incomes.

Chart 3: Drop in GDP partly reflects a steep fall in measured output of public services. This has no impact on household incomes



Sources: ONS and Bank calculations.

Chart 4: Households protected from fall in national income



Sources: ONS and Bank calculations.

The third – and almost certainly the most important – factor is the range of furlough schemes put in place this year. These have introduced an even bigger wedge between GDP and household incomes, shielding consumers from the huge drop in national income. This has been done only at the expense of a material rise in public-sector debt, something that will have to be paid for over time. But as far as households are concerned, the schemes have succeeded in smoothing out that cost, helping to sustain both incomes and consumer spending in the meantime.

Let me now turn to another important contributor to the anomaly in Chart 1, and the main focus today – the extent to which consumers have responded to the pandemic by substituting some forms of spending for others. I'll begin with what some simple economics has to say about what you might expect to happen, depending on the circumstances, and then look at a bit of evidence.

The direct impact of the pandemic on aggregate consumption

Imagine the government suddenly slaps a big excise duty on something people buy. Let's suppose – just for the sake of (wholly unrealistic) argument – it's restaurant meals. Imagine too, just to make the parallel a little closer, that the government compensates you for the cost of this (it redistributes the proceeds of the tax). So you don't lose out in terms of income. However, going out to eat still costs far more than it did, relative to other stuff. So you'll probably be inclined to go less often.

But what do you do with the money you would've spent? Do you simply switch it to other things, leaving your aggregate consumption unchanged? Or do you reduce your total spending as well, and instead save a little more?

At least for the idealised consumer of simple economic theory, two things turn out to matter for this decision. One is the extent to which there are other, close substitutes for the taxed activity. If you liked cooking at home, or indeed anything else, almost as much as going to a restaurant you'd be more inclined to switch your spending immediately to these other things. Your total consumption wouldn't change much. In general, a hike in a particular sales tax will have smaller effects on aggregate spending, and result instead in more expenditure-switching, if it's levied on something for which there are similar (and untaxed) alternatives.

The other important factor is how long you expect the policy to last. If he thinks the tax hike is permanent, our idealised consumer of economic theory would substitute his spending straightaway towards other goods or services, regardless of whether they're close substitutes. If not – if he expects the government to come to its senses at some point, and reverse the policy – he'll be more inclined to wait until that happens. Sensing his options will improve in future, and rather than switching spending to some other, imperfect substitute today, he'll instead want to shift it to tomorrow. Spending on these other things would probably still go up, but by less than the decline in the taxed activity. Aggregate consumption would therefore decline.

The way to understand this is to recognise that the amount you save involves a decision about how to allocate your spending over time. And for that decision to change, the future has to look different in some way from the present. Unless it affects your income, a permanent change in relative prices won't do much to overall spending. Temporary changes, which shift the relative attractiveness of spending at different points in time, can have much more powerful effects.

The cut in VAT in late 2008, for example, enacted at the height of the financial crisis, was by design in place for only a year. It improved households' real incomes and therefore their capacity to spend. But the more important thing was that it was an explicitly temporary policy. If you know something's going to cost more in future there's an added incentive to buy it now. The same is true of the various policies designed to boost spending in 2020: the VAT cut for the hospitality sector, the "Eat Out to Help Out" scheme and the reduction in Stamp Duty. It's doubtful any of these would have had the same impact had they been seen as permanent changes.

I've gone through this because I think there's a clear parallel with the pandemic and its effects on consumer demand. A jump in the prevalence of the illness acts like a big rise in the cost of particular things, activities that expose you to the risk of infection. The cost isn't a monetary one (which is why a better analogy is a compensated tax hike). But it very clearly deters spending on riskier activities. By the same token, it also increases the incentives either to switch spending immediately to other, less risky things, or to defer it to the future (i.e. to cut aggregate consumption and save more).

And the factors that determine the balance of these responses are the same as for the hypothetical tax hike. If either there are close substitutes for the riskier activities, or if you expect the jump in infection rates to persist for an extremely long time, there will in theory be more limited effects on aggregate spending and a greater degree of switching across categories². Employing the economists' jargon, substitution from the risky activities will be more "intra-temporal" than "inter-temporal". If, on the other hand, people expect the rise in infection rates (and the associated restrictions) to subside relatively quickly – and if there are also no close substitutes for the sorts of spending that expose people to infection risk – one would expect a more marked effect on aggregate consumption, and less expenditure switching.

Charts 5(a)-(c) depict some simulations to get these points across. The model they're based on is extremely simple and stylised. In particular, it assumes that incomes, current and expected, are held fixed, and that there's no change in the degree of uncertainty about the economic outlook. The only thing that changes is the path, actual and expected, of the pandemic.

This is very obviously unrealistic. Household incomes have clearly been affected by the crisis. And we know from long experience that when uncertainty about people's future incomes goes up, saving rates tend to rise as well.

But as we saw earlier, from Chart 4, household incomes have been significantly protected by the furlough schemes, at least as compared with the huge dip in national income. (While this isn't a free lunch – higher government debt has to be paid for eventually – it does spread out the cost of that dip.) So the assumption

² In a more general model than the one used here Krueger et al (2020) also point out that the impact of the pandemic on aggregate spending is lower the closer the degree of substitutability between risky and non-risky consumption. Charalampidis and Guillochon (2020) and Guerrieri et al (2020) also consider Covid in multi-sector models. See Cantore et al (2020) for a brief introduction to the epi-macro literature.

that incomes have been relatively stable may not be as unrealistic as it might have been. It's in any case useful, I think, to try to separate the direct effects of the pandemic on spending from its second-round impact, via incomes and economic risk.

It's these direct effects that these simulations attempt to represent. Chart 5(a) plots two stylised paths for infection risk, one of which lasts a lot longer than the other (blue versus red). Our hypothetical consumer understands these dynamics and reacts accordingly in Chart 5(b). Consumption of the risky good, which at the start is assumed to account for around a quarter of the total, follows (inversely) the path of the infection. Spending on non-risky things moves in the opposite direction – there's a degree of expenditure switching in both cases – but it's more significant, and the initial and corresponding fall in aggregate spending less marked, in the blue case, when the virus is expected to persist for longer. Thus the response of consumer spending in aggregate is either drawn out but relatively limited (blue) or sizeable but shorter-lived (red). Note that, in both cases, aggregate consumption is slightly higher once the infection risk has gone because, having been saving more, the consumer has greater net wealth by that point.

Chart 5a: Two stylised paths for the pandemic, one more persistent than the other

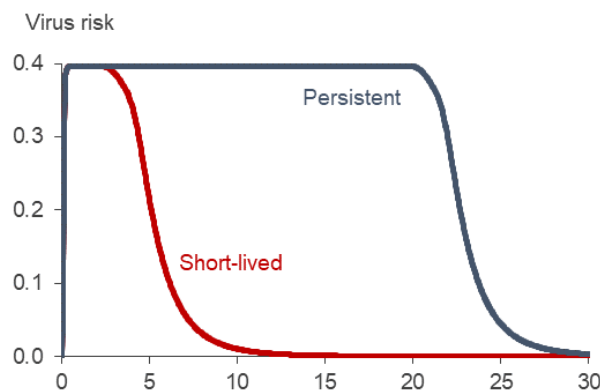


Chart 5b: Sharper response when infection rates less persistent

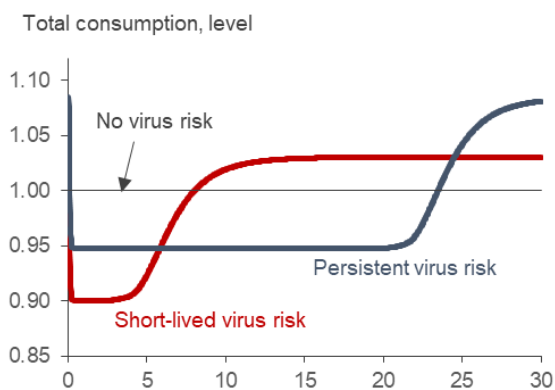



Chart 5c: Response smaller when there are close substitutes for risky spending



Finally, Chart 5(c) demonstrates what happens to aggregate spending when the two categories of spending are assumed to be closer substitutes³. Both lines use the shorter-lived spike in infection rates (the red line in


³ The degree of substitutability between two goods corresponds to the sensitivity of their relative demand to a change in their relative price (i.e. the slope of the demand curve). Chart 5(b) and the red line in Chart 5(c) assume that, in normal circumstances, a 1% rise in its relative price would reduce demand for the risky good by 1½%. The grey line in Chart 5(c) doubles this assumed response to -3%.

5(a)). When it's easier to find a reasonable substitute for the risky spending then, unsurprisingly, that's what the consumer does. The drop in aggregate spending is also smaller in this case.

How do these points square against the real-world experience? It's hard to test them definitively, with less than a year's experience to go on. It's also quite possible that some of these parameters have changed through time. During the first lockdown, for example, there may have been a greater share of consumption that was effectively considered "risky" than in the more recent episodes⁴. Consumers may also have learned over time how to adapt to these restrictions a little more. But I think there are still a few things in the empirical data that are worth picking out.


First, I think there clearly was a fair amount of expenditure switching through this year. A pound less spent on risky activities like going to restaurants or concerts has not meant a pound less in aggregate. This is most obvious in cases where there are relatively close substitutes for the things people have been unable to do. For example, the surge in retail spending we saw at the start of this talk has occurred despite a big drop in sales at physical shops. That has been more than made up for by a big rise in online purchases (Chart 6). Similarly, the inability to go out in people's free time has probably contributed to a strong rise in spending on audio-visual equipment, as they instead try to improve their leisure time at home. More generally, and as we saw at the start, goods spending has been much stronger than spending on services – and also stronger than the average rate of growth over the past and that of household income.

Chart 6: Online sales jumped during lockdown



The share of online retail sales. Source: ONS.

Chart 7: Saving rose and fell sharply during and after the spring lockdown



Sources: Bank of England, ONS, The Investment Association, and Bank calculations. The blue line includes flows into household deposits with banks, National Savings and Investment accounts (NS&I), and retail funds.

⁴ I'm not going to distinguish here between consumers' own inclinations and any additional impact of mandated lockdowns. That the lockdowns have been necessary is clear. While they obviously have an incentive to avoid infection themselves individuals won't always take fully into account the risk they pose to others, particularly when so many people can apparently be infectious but free of symptoms. In economists' language there's a clear "externality" that obligatory lockdowns are seeking to address. For the same reason, it seems very likely that the lockdowns will have had an additional effect on spending, above and beyond those that would've resulted from consumers' own natural caution (although as my colleague Jan Vlieghe recently explained (Vlieghe (2020)), it's empirically difficult to separate these things in the data because, in most countries, they've coincided with each other, both of them responding to swings in infection rates). In the current context you can think of this as a widening in the basket of activities that carry infection risk – and perhaps, more fundamentally, as the result of a model in which individual utility is replaced by some measure of social wellbeing.

Second, however, there are many affected areas of consumption for which there are no perfect (or even very good) substitutes. And as far as the dynamics are concerned, movements in infection rates have been extremely rapid – in Chart 5(b) we’re definitely in the red corner rather than the blue. This may have contributed to the drop in aggregate spending we saw in the second quarter of last year. If you anticipate that the worst will soon be over, and that the conditions for consumption will soon improve, you’re much more likely to delay some of it. Switching there may have been, but aggregate spending still declined and the saving rate still jumped (Chart 7).

Third, and equally, the rapid dynamics are evident in both directions. Total consumption fell sharply in the spring, spending on risky activities even more so. But they both came back pretty smartly in the summer – much faster than in any normal cyclical upswing – as infection rates declined and lockdown was lifted. The jump in the saving rate in the second quarter was easily the largest rise on record in the (near) 60-year quarterly series. The drop in Q3 was also the largest fall on record. On average, in the UK data, the “half-life” of deviations in the saving rate – the time it takes to retrace half of any change – is over a year. On this occasion it happened in a matter of weeks.

The idea that consumer spending would come back quite quickly once infection rates had begun significantly to decline – faster, in particular, than in the “average” economic recovery – was something built in to the MPC’s “scenario” last May, early on in the pandemic. At the time, there was quite a bit of scepticism about this and many judged it to be too optimistic. For my part I think there was every reason to expect a relatively sharp recovery – slower, perhaps than the precipitous drop but nonetheless faster than in a typical economic upswing. In the event spending came back even faster during the summer than we’d expected in May. The same was true in other countries. And the interesting thing about these simple simulations (to the extent one attaches any credence to them) is that they suggest that the initial downturn in aggregate spending may have been as sharp as it was precisely because it was likely to be short-lived.

The role of the furlough schemes was obviously critical here. One of the reasons “normal” recoveries in consumption are more protracted is because drops in output feed through to household incomes, via hits to employment and wage growth. This feedback slows down the dynamics of the average cycle. The schemes helped to interrupt that loop.

By the same token, none of this means that, over the medium term, economic risk won’t continue to affect consumers’ appetite for spending, as it’s always tended to in the past. If rates of unemployment were to remain relatively high even two or three years from now, it’s likely the saving rate will also be above average.

But I do think it’s these more conventional factors – economic risk rather than pandemic risk – that we should have in mind when thinking about what might happen to saving and consumption over the medium term.

Indeed I think this would be true even without taking into account the likely improvements in health treatments (the new vaccines, for example).

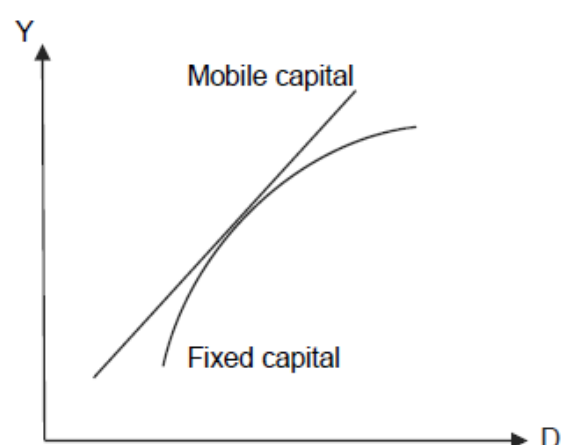
The consequences for supply and costs

A market economy is a dynamic place. Adapting to economic conditions, meeting changes in consumer tastes and adopting new production methods – this is what it's meant to do. And underneath the macro-economic surface, there is always a significant degree of churn of this sort. Even in normal times jobs are destroyed and created all the time, buildings re-purposed and new techniques learned.

Most of the time, as macro-economists, we can safely ignore these things (we certainly tend to). We assume that aggregate quantities are determined only by other aggregates – potential output depends only on total capital and labour, for example, and inflation only on some whole-economy version of spare capacity.

As long as economic resources are mobile this is probably a reasonable thing to do. When thinking about what determines supply capacity at the level of an individual sector we generally think the process exhibits what are called “constant returns to scale”. If you double all the resources involved you double potential output. So in Chart 8a the response of output to demand looks like a straight line. And the nice thing about this is that, if we add up demand across a lot of sectors it doesn't matter what the distribution of that demand looks like. Demand in individual sectors can move anywhere along that straight line and all that will matter for total output is the total amount of spending⁵.

Chart 8a: In the short run there are diminishing returns to scale



But economic resources aren't perfectly mobile, not in the short run at least. Even if people can move relatively quickly from one job to another, it takes time to repurpose a building and to retrain an employee. And when some of those capital resources are slow to move supply looks more like the curved line in Chart 8a: there are “diminishing returns to scale”. This implies that shifts in demand can, even for a given aggregate, lower aggregate supply.

Chart 8b is meant to get this across. Suppose two sectors start with the same level of demand D and that they each have the right level of resources. Then demand shifts from one to the other (to D_L and D_H respectively). When everything can move in response, capital as well as labour – we're on the straight line –

⁵ Samuelson and Swamy (1974).

aggregate output is unchanged. But when we're on the curved line, and averaging across the two sectors, total attainable output is lower. At least until capital is able to react, and move from the low-demand to the high-demand sector, resources are misallocated.

The counterpart for costs is shown in Chart 8c. With diminishing returns marginal costs increase along with output – the curved line gets steeper as you move to the right. So the rise in costs for a high-demand sector is bigger than the drop for a low demand sector. At least in the short run, therefore, big shifts in demand can raise costs across the economy as a whole, even for a given aggregate.

Chart 8b: Potential output could be affected by large and rapid rise in the dispersion of demand





Chart 8c: Dispersion in demand can raise short-run average costs



I shouldn't exaggerate these effects. They don't negate the point that shortfalls in aggregate demand tend to drag on costs and inflation. In reality, in fact, the shift could itself create such a shortfall. Capital may be slower to move but even the labour market doesn't clear immediately and some of those losing their jobs in the low-demand sector would probably spend some time being unemployed before finding new jobs in the high-demand sector.

Note too that the demand shifts would have to be quite big to have any material or noticeable effect. Because the curved and straight lines have the same slope to start with, small shifts don't do that much to average costs in the economy^{6,7}. Indeed, in the real world, recognising that demand can move around unpredictably, firms generally operate with a margin of spare capacity, precisely so they're able to meet any sudden rise in orders without incurring undue extra costs. This serves to flatten out the cost curve a bit.

⁶ Students of economics will recognise this as an example of something called the "envelope theorem". Around an optimum – in this case an optimal amount of capital – changes in the environment have only second-order effects on decisions.


⁷ Clearly the effects could be higher if the economy isn't starting from an allocative optimum to begin with. There is quite a bit of work in the economics literature exploring the effect of allocative inefficiencies on aggregate productivity, even in standard business cycles. See, for example, Davis and Haltiwanger (1998), Chari et al. (2007) and Osotimehin (2019).

But on occasion sizeable shifts do happen. And at least until resources can move in response, they may at the margin add to aggregate costs and inflation.

I wondered about this effect first in the years after the financial crisis, and specifically whether it might help to explain why inflation had been less weak than one might have expected, given the large decline in GDP⁸. I thought at the time that one indicator of such a demand-driven misallocation might be a noticeable sectoral divergence in prices, mirroring that in demand and marginal costs. Chart 9 plots the sectoral dispersion in inflation rates, across broad components of the CPI, over the past 15 years or so. You can see that there was a spike in the series following the financial crisis and another last year.


This isn't a particularly clean indicator. The effect I'm discussing here involves cross-sectoral shifts in demand specifically. But relative prices can change for other reasons too (sector-specific changes in costs or productivity, for example), in a way that wouldn't reflect any sort of misallocation of resources. In an effort to isolate areas where demand looks to have been the dominant influence, Chart 10 plots changes in inflation and sales through the course of last year, each dot indicating a different part of consumption, but only for those sectors where the two things went in the same direction⁹. Tentatively it suggests that, in areas where demand was relatively strong, prices went up by a little more than the equivalent declines in low demand areas. This is the kind of "convexity" that Chart 8c was trying to get across.

Chart 9: Dispersion in CPI inflation has risen



Sources: ONS and Bank calculations.

Chart 10: Some evidence of convexity



Sources: ONS and Bank calculations.


It's also true that, so far at least, the huge drops in output in the first half of 2020 have not produced the declines in underlying inflation that forecasters expected in the spring. Chart 11 plots an estimate of "core" CPI inflation in the UK, stripping out the more volatile food and energy components but adding back in what we think were the direct effects of the policies designed to assist the hospitality sector (the "Eat Out to Help Out" scheme and the cut in VAT). In November, this measure of prices was 1.4% higher than a year earlier, close to where it was at the end of 2019.

⁸ Broadbent (2012).

⁹ Supply shocks would tend to move output and prices in opposite directions.

I wouldn't describe this as anything more than suggestive. The international evidence is more mixed (as in the UK, core inflation has been relatively stable in the United States, but it's fallen notably in the Euro area). And even if this mechanism has been at work, it's not clear it has important implications for inflationary pressure over the medium term, which is the focus for monetary policy.

Chart 11: Drop in core inflation moderate so far



Sources: ONS and Bank calculations.

It's reasonable to expect the degree of dispersion in demand and spare capacity to narrow as and when aggregate spending recovers (those areas most affected by the pandemic are likely to see stronger growth afterwards). To the extent shifts in consumption do persist – some have speculated, for example, that a higher proportion of people will work from home in future, with possible implications for the pattern of spending over the longer term – one would expect the economy to adapt to those over time.

But alongside other possibilities that the MPC described in the August *Monetary Policy Report*, the significant dispersion in demand might be one other thing that's effectively served to flatten the price Phillips curve (i.e. to temper the response of inflation to shifts in aggregate spending), if only in the shorter term¹⁰.

Conclusion

The pandemic, and the waves of the restrictions that have accompanied it, have had highly selective effects. They've hit some areas of spending much harder than others.

The decline in aggregate consumer spending would almost certainly have been more protracted and more widespread – affecting not just activities that involve infection risk but those that do not – but for the furlough schemes. These have helped shield household incomes from the big drop in national income last year.

But that's not the only thing going on. Spending on some of these “non-risky” things have actually risen faster than incomes – in some cases faster than they have for many years. So it looks as though consumers used some of the money they would have spent on riskier activities – going to restaurants and so forth – to buy other things instead.

¹⁰ Chapter 4 of that Report went through a number of reasons why, at least in the short term, the response of inflation to demand might be smaller than usual under these circumstances. Services prices tend to be “stickier” than average, so a downturn centred on consumer services may be less disinflationary in the first instance. If orders take time to respond, a cut in prices will initially reduce a firm's profits. When balance sheets and cashflow are under pressure, firms may therefore be more reluctant to make such cuts.

This makes sense. A pound less spent on one thing needn't mean a pound less in total. Indeed under some (admittedly extreme) conditions you'd expect substitution of this sort to be one-for-one. If people are entirely indifferent between different sorts of consumption – if they're just as happy with “non-risky” as “risky” stuff – aggregate spending wouldn't fall at all (for given income). What they saved on one would be spent in its entirety on the other. The same is true if a rise in infection rates is expected to be permanent. If the future looks exactly like the present there's no reason for saving – or therefore total consumption – to change.

These conditions are obviously unrealistic. Some things are good substitutes for risky activities during the pandemic. Perhaps it's not much harder to buy things online than in a shop. At a stretch, watching a film on a screen at home can replace a trip to the cinema. But going out with family or friends to a pub or a restaurant, listening to live music or watching a play – for many of the things we've been denied for much of the past year, there is no good proxy. As for the dynamics of the pandemic, one thing they're not is permanent or even slow-moving. Infection rates went up very rapidly in the spring and, once lockdown was in place, they fell almost as rapidly in the early summer. With the promise of looser restrictions as that happened, it would make sense for people to hold off on some spending until the day itself arrived.

Nevertheless, it's clear that a fair degree of expenditure switching took place: some areas of consumption fared unusually well last year, even as others declined extremely sharply. Because it's hard for economic resources to respond to large and rapid shifts in demand this may have softened somewhat the immediate impact of the downturn, enormous as it was, on inflation.

Whether that has any implications for the medium term is less clear. The dispersion in demand and spare capacity is likely to narrow during the recovery phase. Over time, if some of these shifts in demand prove more persistent, the economy is likely to adapt to them and their effect on aggregate supply and costs to diminish.

More generally, the rapid dynamics of the virus, combined with the vast scale of the furlough schemes (which act like extremely large automatic stabilisers for employment), make it more difficult than usual to infer what short-term gyrations in GDP mean for inflationary pressure further ahead. The lockdowns mean that, on a quarterly basis, GDP is likely to have fallen in the fourth quarter of 2020 and to do so again in 2021Q1. This will no doubt prompt headlines about a “double-dip recession”. But this is quite unlike any normal economic cycle. Its size, shape and speed are all highly unusual. As far as medium-term inflation is concerned, what matters more than these gyrations in output per se – very large though they are – are developments over time in the labour market. The rate of unemployment is the best single measure we have of economic slack. It also affects households' confidence about their income and is therefore correlated with rates of saving. And it's because, unfortunately, we expect unemployment to rise once the furlough schemes are wound down (the MPC's central forecast in November was for a peak rate of almost 8%, despite a sharp projected recovery in output as health risks fade) that the appropriate response has been to ease policy significantly. The Committee has also said that it would need firm evidence of a significant narrowing in spare capacity,

and of a sustainable return of inflation to the 2% target, before considering whether to withdraw any of this stimulus.

Come what may, the MPC will continue to set policy in order to achieve the 2% inflation target, in line with its remit, over the medium term.

Thank you.

References

- Broadbent, B., 2012.** 'Productivity and the allocation of resources'. Speech given at the Durham Business School. Available at <https://www.bankofengland.co.uk/speech/2012/productivity-and-the-allocation-of-resources>.
- Cantore, C., Di Pace, F., Masolo, R. M., Miranda-Agrippino, S., and Turrell, A., 2020.** 'Covid-19 briefing: epi-macro 101'. *Bank Underground*, available at <https://bankunderground.co.uk/2020/08/07/covid-19-briefing-epi-macro-101/>.
- Charalampidis, N. and Guillochon, J., 2020.** 'The COVID-19 pandemic and the consumption of nondurables and services'. Available at SSRN: <https://ssrn.com/abstract=3685336>.
- Chari, V. V., Kehoe, P. J. and McGrattan E. R., 2007.** 'Business cycle accounting'. *Econometrica*, Vol. 75, pp. 781-836.
- Davis, S. J. and Haltiwanger, J., 1998.** 'Measuring gross worker and job flows'. In: Labour statistics measurement issues. *University of Chicago Press*, pp. 77-122.
- Guerrieri, V., Lorenzoni, G., Straub, L. and Werning, I., 2020.** 'Macroeconomic implications of COVID-19: can negative supply shocks cause demand shortages?'. *NBER working paper* No. 26918.
- Krueger, D., Uhlig, H. and Xie, T., 2020.** 'Macroeconomic dynamics and reallocation in an epidemic: evaluating the "Swedish solution"', *NBER Working Paper* No. 27047.
- Osotimehin, S., 2019.** 'Aggregate productivity and the allocation of resources over the business cycle', *Review of Economic Dynamics*, Vol. 32, pp. 180-205.
- Samuelson, P. A. and Swamy, S., 1974.** 'Invariant Economic Index Numbers and Canonical Duality: Survey and Synthesis'. *The American Economic Review*, Vol. 64, pp. 566-593.
- Vlieghe, G., 2020.** 'Assessing the health of the economy'. Speech available at <https://www.bankofengland.co.uk/-/media/boe/files/speech/2020/assessing-the-health-of-the-economy-speech-by-gertjan-vlieghe.pdf>.